arXiv:hep-ex/9904025v1 23 Apr 1999

Estimation of Upper Limits Using a Poisson Statistic

Ilya Narsky [1]

Physics Department, Southern Methodist University, Dallas, TX, 75275-0175, USA

PACS: 02.50.-r, 02.50.Cw, 02.50.Ng.

Abstract

Bayesian, classical, and extended maximum likelihood approaches to estimation of upper limits in experiments with small numbers of signal events are surveyed. The discussion covers only experiments whose outcomes are well described by a Poisson statistic. A new approach, based on the statistical significance of a signal rather than on the number of events in the signal region, is proposed. A toy model and an example of a recent search for the lepton number violating decay $\tau \to \mu\gamma$ are used to illustrate application of the discussed techniques.

1 Introduction

Searches for rare signals often fail to detect a signal of sufficient statistical significance and thus face the need to set an upper limit on the signal rate. Unfortunately, there is no standard prescription for setting such limits, and a number of techniques have been employed in the past to meet this challenge. This problem is particularly important for the particle physics community. Upper limits on rare and forbidden decays can provide valuable constraints on physics beyond the Standard Model of Electroweak Interactions.

This paper is stimulated by the observation that often authors do not pay enough attention to the choice of a procedure for upper limit estimation. Furthermore, many experimentalists do not provide a sufficient description of the procedure used in their analysis technique, assuming that the reader is smart enough to figure out the details. An example given in Section 2.2 shows

[1] Tel.: 1 214 768 1299; Fax: 1 214 768 4095; E-mail: narsky@mail.physics.smu.edu.

that, in some situations, the choice of a procedure can change the value of an upper limit *by an order of magnitude*, which suggests that this problem should not be taken lightly.

In Section 2 the most popular approaches to upper limit estimation are surveyed and criticism is provided, where applicable. In Section 3.2 definitions of the statistical significance of a signal are discussed and a new technique, which estimates upper limits using these definitions, is proposed. Unlike commonly used procedures, this approach does not just count the number of events in the signal region, but takes into account statistical fluctuations of the background as well. In Section 4 various techniques are compared using a toy model. In Section 5 an example of a recent search for the lepton number violating decay $\tau \to \mu\gamma$ [25], performed by the CLEO collaboration, is given.

The emphasis is made on estimation of upper limits, though the discussion can be easily generalized to include construction of confidence intervals whose lower bound is not constrained to zero.

There is no discussion of systematic effects in this paper. Throughout the paper the number of observed events, the expected background rate and, if applicable, the coordinates of observed events are assumed to be measured with perfect accuracy.

2 Commonly Used Techniques for Upper Limit Estimation

2.1 Bayesian Approach

In a Bayesian approach one has to assume a *prior* probability density function (pdf) of an unknown parameter and then perform an experiment to update the prior distribution. A prior pdf reflects knowledge that is available to an experimentalist *before* an experiment is performed. The updated prior is called the *posterior* pdf and is used to draw inference on the unknown parameter. This updating is done with the use of Bayes' Rule. For the moment let us ignore the issue of background, i.e., let us assume that the background rate is measured very accurately and thus can be treated as a known constant. Then the only unknown parameter is the signal rate s. Bayes' Rule gives:

$$\pi(s|n) = \frac{f(n|s)\pi(s)}{\int_0^\infty f(n|s)\pi(s)ds} \ , \tag{1}$$

where n represents the number of observed events, $f(n|s)$ is the conditional probability to observe n events, given the signal rate s, $\pi(s)$ is the prior pdf,

and $\pi(s|n)$ is the conditional posterior pdf. Now, for any given confidence level, one can compute a Bayesian confidence interval for the signal rate s. For upper limit estimation the natural choice of the Bayesian confidence interval is of the form $(0, s_0)$. Here, s_0 denotes an upper limit and can be found from the equation:

$$1 - \alpha = \int_0^{s_0} \pi(s|n)ds \ . \tag{2}$$

The confidence level is denoted by $(1-\alpha)$ following the conventional statistical notation. A nice feature of the Bayesian approach is that the zero value of an upper limit s_0 always corresponds to the zero value for the confidence level $(1 - \alpha)$. As you will see below, this is not necessarily true for the classical approach.

The most important step here is to define a prior distribution of the parameter. Naturally, this is the step which brings most of controversy (and sometimes confusion) into Bayesian methods. Many statistical textbooks treat a prior pdf as a purely subjective assumption which is based on experimenter's belief. At the same time there are authors [3–6] who advocate the objectivity of prior assumptions. In particular, Jaynes [5] stated the "basic desideratum" of the objective approach as follows: *in two problems where we have the same prior information, we should assign the same prior probabilities.* These authors have offered a number of mathematical procedures that can be used to convert prior knowledge into an exact formula for the prior distribution. Another important question is whether one should assume an *informative* prior, i.e., a prior which incorporates results of previous experiments, or a *non-informative* prior, i.e., a prior which claims total ignorance. Naively, it would seem unnatural to use non-informative priors since the power of the Bayesian approach comes from the fact that we update our prior knowledge rather than start every analysis from scratch. In particle physics the major objection against informative priors is based on the following argument: if we assume a prior which incorporates results of earlier experiments, then our experiment will not be independent of those and thus we will not be able to combine our results with the results of previous experiments just by taking a weighted average. So far, particle physicists have been largely ignoring informative priors. Thus, the discussion below covers only those Bayesian methods that assume a non-informative prior pdf for the positive parameter of a Poisson distribution.

In the absence of background the conditional pdf $f(n|s)$ is given by:

$$f(n|s) = e^{-s}\frac{s^n}{n!} \ . \tag{3}$$

Bayes and Laplace [1,2], who pioneered statistical research employing Bayesian methods, stated that the non-informative prior for any parameter must be flat. This conclusion was not based on any strict mathematical argument, it was merely a product of their intuition. Modern advocates [8–11] of this approach do not offer any mathematical explanation either, they just consider a flat prior pdf as the most natural choice one can make. As natural as this assumption may seem, there are several objections. The most obvious argument is that if one can assume a flat distribution of an unknown parameter, then one can also assume a flat distribution for any function of this parameter, and these two assumptions are clearly not identical. If an experiment is measuring the mean of the Poisson statistic (3), then one can argue that the most natural candidate for the unknown parameter is the signal rate s, and we have no physical reasons to consider any functions of this parameter except the parameter itself. In other experiments the situation may be not so simple. For example, if an experiment is measuring a neutrino mass, then typically the measured quantity is not the neutrino mass itself, but the neutrino mass squared. In this situation it's not clear whether one should choose mass or mass squared as a candidate for the flat distribution. Jeffreys [3] resolved this problem by introducing an *invariate* prior pdf $1/\theta$ which, he stated, was a valid non-informative prior for all problems where the unknown parameter θ could vary from 0 to $+\infty$. His choice was mostly motivated by the fact that $d\theta/\theta \propto d\theta^n/\theta^n$, i.e., the pdf of θ stays invariant under any power transformation. All non-power functions of θ were rejected by Jeffreys as non-physical under the assumption that the parameter θ is a dimensional quantity. This argument was put on a more rigorous mathematical basis by Jaynes [5] who stipulated that a prior pdf has to stay invariant under any symmetry transformation [2] that does not change the physics of an experiment. His conclusion was similar to that of Jeffreys, namely that the non-informative prior for a Poisson statistic (3) has to be proportional to $1/s$. An alternative approach was developed by Box and Tiao [7] who introduced the notion of a *data-translated* likelihood. In their approach a prior pdf is non-informative if the location, but not the shape, of the corresponding posterior likelihood [3] is determined by the unknown parameter. Thus, the location of the posterior likelihood is completely

[2] Strictly speaking, Jaynes' argument is not applicable to the Poisson distribution (3) with a non-dimensional mean rate s. Originally Jaynes' non-informative prior was derived for the Poisson pdf $f(n|s) = e^{-st}(st)^n/n!$, where n is the number of counts, t is the counting time, and s is the signal rate per unit time. However, the requirement of dimensionality seems to be somewhat arbitrary. For example, in the above formula t can be the amount of statistics accumulated in the experiment and s can be the signal rate normalized to this amount of statistics. A prior pdf has to stay invariant under the transformation $s' = qs$, $t' = t/q$, i.e., Jaynes' argument is applied.

[3] As usually, a likelihood function is defined by swapping argument and parameter in the expression of the corresponding pdf.

defined by data, while its shape had been determined before the data were seen (hence the name "data-translated" likelihood). It may not be possible to construct a data-translated likelihood for every distribution. In this case one can use Taylor expansion to construct an approximate data-translated likelihood. In particular, a non-informative prior which produces an approximate data-translated likelihood for the Poisson statistic (3) is given by $1/\sqrt{s}$.

In the presence of background, the Poisson pdf (3) has to be modified to account for the non-zero background rate b:

$$f(n|s) = e^{-(s+b)}\frac{(s+b)^n}{n!} \ . \tag{4}$$

If the background rate is accurately measured, it can be treated as a known constant. In this case an argument similar to that of Ref. [5] gives $1/(s+b)$ and an argument similar to that of Ref. [7] gives $1/\sqrt{s+b}$ for the non-informative prior. In general, for the prior pdf

$$\pi(s) \propto \frac{1}{(s+b)^m} \ ; \quad 0 \le m \le 1 \ ; \tag{5}$$

the posterior distribution is given by

$$\pi(s|n) = \frac{e^{-(s+b)}(s+b)^{n-m}}{\Gamma(n-m+1,b)} \ , \tag{6}$$

where

$$\Gamma(p,\mu) = \int_{\mu}^{\infty} s^{p-1}e^{-s}ds; \quad p>0; \quad \mu>0; \tag{7}$$

is an incomplete gamma-function. Substituting the posterior pdf (6) into Eqn. (2), we obtain:

$$1-\alpha = 1 - \frac{\Gamma(n-m+1, s_0+b)}{\Gamma(n-m+1,b)} \ . \tag{8}$$

A gamma-function is not well-defined if its first argument is equal to zero or a negative integer. Thus, at $n=0$ (no events observed) and $m=1$ (the $1/(s+b)$ prior) Eqn. (8) fails to find an upper limit. Under the flat prior $m=0$, Eqn. (8) turns into the formula

$$1-\alpha = 1 - \frac{\sum_{k=0}^{n} e^{-(s_0+b)}\frac{(s_0+b)^k}{k!}}{\sum_{k=0}^{n} e^{-b}\frac{b^k}{k!}} \ , \tag{9}$$

which was adopted by Particle Data Group [13]. Here and below, 0^0 is by definition equal to one.

If the background rate is measured with large uncertainty, then the situation is more complicated. The usual approach is to modify the Poisson statistic (4) as:

$$f(n|s) = \int_0^{+\infty} e^{-(s+b)} \frac{(s+b)^n}{n!} f(b) db \ , \tag{10}$$

where $f(b)$ is the measured or predicted pdf of the background rate b. The argument of the previous paragraph, leading to the $1/(s+b)$ and $1/\sqrt{s+b}$ expressions for the non-informative prior, is not necessarily valid in this case. A more consistent Bayesian approach would be to derive a joint non-informative prior $\pi(s,b)$, thus treating both the signal and background rates as unknown parameters. An example of such derivation is given in Ref. [6]. A discussion of intricacies, that might arise from inclusion of background uncertainty into the upper limit calculation, is beyond the material covered in this paper.

2.2 Classical Approach

The classical (or frequentist) approach is traditionally interpreted in the following way: if a $(1-\alpha)$ classical confidence set is constructed for the unknown parameter, then the probability for this confidence set to cover the true value of the parameter equals $(1-\alpha)$, that is, if an infinite number of identical independent experiments is performed and for each of these a $(1-\alpha)$ confidence set is constructed, then $100(1-\alpha)\%$ of these confidence sets will and $100\alpha\%$ of these confidence sets will not contain the true value of the parameter. Rules for construction of classical intervals were outlined in a famous work [15] by Neyman. Here I treat the terms "classical" and "frequentist" as equivalent. The implication is that the classical approach is inevitably connected to the concept of an identical independent experiment. Sometimes this concept is misunderstood, and are curious attempts to treat certain problems in the "classical" vein while, in fact, this treatment has nothing to do with the frequentist approach. An example of such misunderstanding is shown in this Section.

A confidence set for the unknown signal rate s is typically constructed [16] as a confidence interval $s_1 \le s \le s_2$ satisfying:

$$1-\alpha = \sum_{k=n}^{\infty} f(k|s_1) + \sum_{k=0}^{n} f(k|s_2) \ . \tag{11}$$

For upper limit estimation it is natural to consider intervals $0 \leq s \leq s_0$, where s_0 is the value of an upper limit. Eqn. (11) is thus reduced to:

$$1 - \alpha = 1 - \sum_{k=0}^{n} e^{-(s_0+b)} \frac{(s_0+b)^k}{k!} \ , \tag{12}$$

where $f(k|s_0)$ was replaced by its definition (4).

Eqn. (12) looks similar to the Bayesian formula (9), except that the denominator is now absent. The denominator in Eqn. (9) represents the probability of observing n or less background events in the signal region, and thus, it is always less than 1 except at $b = 0$. Therefore, for the non-zero background rate, i.e., $b \neq 0$, the classical approach (12) always provides a smaller value of an upper limit than the Bayesian approach (9) with a flat prior. The difference becomes significant when the denominator in Eqn. (9) is small, i.e., when the number n of events observed in the signal region is small as compared to the expected background rate b. For example, for the observed number of events in the signal region $n = 3$ and the expected background rate $b = 6.5$ the classical approach (12) gives an upper limit $s_0 = 0.18$ at 90% confidence level, while the Bayesian approach (9) with a flat prior gives $s_0 = 3.39$, i.e., a difference of more than an order of magnitude. Another feature of Eqn. (12) is that now a zero upper limit: $s_0 = 0$, does not give you a zero confidence level. When the expected background rate b is large, one can achieve a situation when $1 - \alpha < 1 - \sum_{k=0}^{n} e^{-b} b^k/k!$, i.e., formula (12) gives an unphysical negative value of an upper limit. Advocates of the Bayesian method consider this as an undesirable feature of the classical approach. But the failure to set an upper limit is not necessarily a bad feature. It just implies that the observed outcome of an experiment is highly improbable, and one should question the experimental technique that was used to measure the signal and the background.

Zech [12] made an attempt to derive the Bayesian formula (9) using the classical approach. To arrive at Eqn. (9), Ref. [12] postulates that the number of background events cannot exceed the number n of observed events and therefore one has to renormalize the probability (12) according to this constraint. This approach misunderstands the concept of frequentist coverage. Eqn. (12) answers the question: what is the probability of observing n or more events in an *identical independent* experiment? This implies that the number n of events cannot be fixed at any particular value. If you fix n, then you confine yourself to this particular drawing; the concept of an identical independent experiment is therefore inapplicable. In this case a binomial distribution should be used instead of a Poisson pdf.

Confidence intervals are not uniquely defined by Eqn. (11) unless one imposes specific criteria on their construction. In the situation discussed above, the

uniqueness was introduced by requiring that the confidence set must be an interval $(0, s_0)$. One can choose another requirement and obtain a different confidence set. In general, if the problem is reduced to one variable and one unknown parameter, then confidence intervals are obtained via construction of confidence belts [20]. The idea behind this technique is the following. For every assumed value of the signal rate s, one finds an acceptance interval $n_1(s) \leq n \leq n_2(s)$ which satisfies:

$$1 - \alpha = \sum_{n=n_1}^{n_2} f(n|s) \ . \tag{13}$$

Due to the discrete nature of the Poisson distribution, it is usually impossible to find values of n_1 and n_2 that satisfy Eqn. (13). The standard solution is to stay on the conservative side and to search for an acceptance interval that gives at least the required coverage:

$$1 - \alpha \leq \sum_{n=n_1}^{n_2} f(n|s) \ . \tag{14}$$

The obtained functions $n_1(s)$ and $n_2(s)$ define two curves on the s-vs-n plane. Then, for every value of n, one obtains an interval $s_1(n) \leq s \leq s_2(n)$ which lies between these two curves. The obtained interval is a $(1 - \alpha)$ confidence interval for the specific value of n. The order in which values of n are added to the acceptance region (n_1, n_2) for a specific value of s is called an *ordering principle*. Thus, every ordering principle corresponds to a specific set of confidence intervals (s_1, s_2) indexed by the observation variable n. The well-known examples are an old paper [18] by Crow and Gardner and a recent paper [17] by Feldman and Cousins. Ref. [18] minimizes the length of the acceptance interval (n_1, n_2) defined by Eqn. (14). Ref. [17] employs an ordering principle based on likelihood ratios. The latter was adopted by the last release of the Particle Data Group Review [14]. In both approaches an experimentalist does not have to decide whether she/he wants to quote an upper limit or a confidence interval: she/he simply applies the chosen procedure to construct a confidence interval (s_1, s_2). If the lower bound turns out to be strictly equal to zero: $s_1 = 0$, then an upper limit is quoted, and if the lower bound is positive: $s_1 > 0$, then a confidence interval is quoted. This versatile procedure, however, has one subtle problem. The fact that the constructed 90% confidence interval has a non-zero lower bound does not guarantee that a signal of high statistical significance is observed. There is nothing surprising about that. A measured signal rate is usually quoted if the statistical significance of the signal exceeds 3 which corresponds to 99.87% of the area under a Gaussian peak. At the same time, typical confidence levels used to quote upper limits are 90% and 95%. For example, at $b = 1$ and $n = 3$ the 90% confidence interval obtained by the procedure of Ref. [17] equals $(0.10, 6.42)$. At the same time, the statistical

significance σ calculated as $1/\sqrt{2\pi}\int_{\sigma}^{\infty} e^{-x^2/2}dx = \sum_{k=n}^{\infty} e^{-b}b^k/k!$ is only 1.4. Thus, the non-zero lower bound has no clear interpretation: it is a vague indication that a clear signal might be observed in a future experiment. On the contrary, if one chooses to construct a one-sided interval and thus quote an upper limit, this clearly expresses the fact that a signal of sufficient statistical significance was not observed and therefore the signal rate is believed to be consistent with zero.

2.3 Unbinned Extended Maximum Likelihood Fit

Unbinned extended maximum likelihood fits [21,22] have become popular in the past few years. This is an excellent analysis tool for rare signal searches since, unlike a standard maximum likelihood, an extended maximum likelihood correctly incorporates the Poisson error on the number of observed events. Thus, when one observes contributions from two processes (e.g., signal and background), one should use an extended unbinned likelihood:

$$L(s,b) = \frac{e^{-(s+b)}}{N!}\prod_{i=1}^{N}(sS_i + bB_i) \tag{15}$$

instead of a standard likelihood:

$$L(f_s) = \prod_{i=1}^{N}(f_s S_i + (1-f_s)B_i) \ .$$

Here S_i and B_i represent the signal and background spatial pdf's, respectively. N is the total number of events observed in the signal region and in the vicinity, and s and b are the signal and background rates, respectively.

It has also become a common practice to extract an upper limit value by integrating a likelihood function:

$$1-\alpha = \frac{\int_0^{s_0} L(s)ds}{\int_0^{\infty} L(s)ds} \ , \tag{16}$$

where s_0, as usually, denotes the value of an upper limit. The likelihood $L(s)$ is typically obtained by integrating the two-dimensional likelihood $L(s,b)$ over the parameter space b:

$$L(s) = \int_0^{\infty} L(s,b)db \ . \tag{17}$$

This technique is applied, for instance, to estimate upper limits in rare B decay studies [23,24]. It is important to realize that the integration of the likelihood implicitly uses the Bayesian approach. As it was pointed out by Cousins [19], *in particle physics, the prior is almost always taken uniform (where non-zero), although this assumption goes unemphasized by those who merely report that they "integrated the likelihood function".*

To understand this comment, one should recall the definition of a likelihood function:

$$L(s,b) = f(\vec{x}|s,b) \ , \tag{18}$$

where $\vec{x} = \{N, x_1, x_2, ..., x_N\}$ is a set of observables, the quantity x_i ($i = 1, ..., N$) is the coordinate of the ith event, the symbol f denotes the corresponding pdf, and a vertical line $|$ always implies conditional distribution. Applying Bayes' Theorem, one obtains:

$$L(s,b) = \frac{f(s,b|\vec{x})f(\vec{x})}{f(s,b)} \ .$$

If signal and background are assumed independent [4], their joint pdf must factorize: $f(s,b) = f(s)f(b)$. Multiplying both sides of the above equation by $f(b)$ and integrating over b, one obtains:

$$\int_0^\infty L(s,b)f(b)db = \frac{f(s|\vec{x})f(\vec{x})}{f(s)} \ .$$

The quantity on the right side of the equation can be recognized as $L(s)$. Therefore,

$$L(s) = \int_0^\infty L(s,b)f(b)db \ . \tag{19}$$

In the small signal limit $S_i \ll B_i$, the likelihood (15) is expressed as:

$$L(s,b) = \frac{e^{-(s+b)}}{N!} b^N \prod_{i=1}^{N} B_i \ .$$

Due to the fact that the contributions from the signal and background rates s and b decouple, the background pdf $f(b)$ in Eqn. (19) cannot change the ratio

[4] The assumption of their independence seems natural. However, Prosper [6] derived a non-informative joint prior $f(s,b)$, where the signal and background contributions are non-factorizable. In fact, this assumption is not as obvious as it may seem.

$\int_0^{s_0} L(s)ds / \int_0^{\infty} L(s)ds$. But when the approximation $S_i \ll B_i$ is no longer valid, this is not so. Thus, Eqn. (17) treats the pdf $f(b)$ as prior knowledge and assumes that it is flat in the interval $(0, +\infty)$.

Furthermore, the posterior distribution $f(s|\vec{x})$ can be represented as:

$$f(s|\vec{x}) = \frac{L(s)f(s)}{f(\vec{x})} \ ,$$

which is similar to the Bayesian formula (1). Thus, Eqn. (16) implicitly assumes a flat prior $f(s)$ and integrates $f(s|\vec{x}) \propto L(s)$ to extract an upper limit. When the experiment is dominated by background, i.e., $S_i \ll B_i$, the likelihood (17) is proportional to e^{-s} and the extracted value of an upper limit is equal to 2.3 at 90% confidence level. The approximation $S_i \ll B_i$ is equivalent to the situation when no events are observed in the signal region ($n = 0$), but the expected number of background events is non-zero ($b \neq 0$). Under these conditions, the Bayesian method (9) with a flat prior gives the e^{-s} behavior as well. Neither the classical approach, nor the Bayesian methods with alternative priors give the e^{-s} behavior for $n = 0$ and $b \neq 0$; all these techniques give smaller upper limit values. Thus, the integration of the likelihood and the Bayesian method with a flat prior give the most conservative estimates in background-dominated analyses.

One can easily avoid the implicit use of the Bayesian method and set an upper limit using the frequentist definition. It is straightforward to implement this approach by running a Monte Carlo job. First, the observed data $\vec{x}_{obs}$ is fitted to the likelihood function (15) to extract estimates s_{obs} and b_{obs} of the true signal and background rates. Then, for every assumed value of an upper limit s_0, a Monte Carlo sample, consisting of a large number of experiments, is generated. For each experiment the number of signal and background events are generated assuming Poisson distributions:

$$f(s) = e^{-s_0}\frac{s_0^s}{s!}; \quad f(b) = e^{-b_{obs}}\frac{b_{obs}^b}{b!}; \quad N = s + b;$$

and the coordinates of events are generated, based on the spatial pdf's S_i and B_i. Then the outcome of every experiment is fitted to the same likelihood function (15) and the distribution of measured signal rates $f(s_{meas})$ is plotted. The confidence level corresponding to this value of s_0 is estimated as:

$$1 - \alpha = \frac{\int_{s_{obs}}^{\infty} f(s_{meas})ds_{meas}}{\int_0^{\infty} f(s_{meas})ds_{meas}} \ . \tag{20}$$

Thus, a value of s_0 is chosen that gives the required confidence level $(1 - \alpha)$. This agrees with the frequentist approach when an experimentalist draws

inference about an unknown parameter on the basis of observed data only, without making any subjective assumption.

3 Calculation of Upper Limits, Based on Statistical Significance of a Signal

The Bayesian and classical approaches described in Sections 2.1 and 2.2 are based on the conditional probability $f(n|s)$, i.e., the probability to observe n events in the signal region, given the signal rate s. However, the number of events observed in the signal region by itself is not so important, since the usual goal of an experiment is not to count events in the signal region, but to observe a signal of high statistical significance. These approaches do not take into account the fact that the number of background events fluctuates too. For example, the classical method (12) estimates the confidence level as a probability to observe a larger number of events in the signal region than the number observed in the experiment, given the assumed value s of the signal rate. But, at the same time, the number of background events in the sideband can fluctuate up to high values or down to zero. The former implies that no signal is observed, and the latter implies that a very clean signal is observed. The methods based on the pdf $f(n|s)$ take into account both these possibilities and thus do not distinguish between signals of high and low statistical significance.

Therefore, one should replace the conditional pdf $f(n|s)$ with the conditional pdf $f(\sigma|s)$, where σ represents the statistical significance of a signal. Before we proceed to derivation of the related mathematical formalism, we need to discuss various definitions of statistical significance.

3.1 Definition of Statistical Significance

If n events are observed in the signal region and b is the estimated background rate, then the statistical significance σ of a signal is defined by

$$\frac{1}{\sqrt{2\pi}} \int_{\sigma}^{\infty} e^{-x^2/2} dx = \sum_{k=n}^{\infty} e^{-b} \frac{b^k}{k!} \ , \tag{21}$$

that is, it represents the probability of observing n or a larger number of background events in an identical independent experiment. The background rate b can be estimated in a number of ways. If the background rate is estimated *independently* of the data seen, e.g., from a Monte Carlo analysis or

from another data sample that is known to contain no signal events, then it is common to assume that the estimated background rate is proportional to the rate observed in the independent experiment:

$$b = \zeta_{ind} n_{ind} \ , \tag{22}$$

where n_{ind} is the number of events observed in the independent experiment, ζ_{ind} is the corresponding scale factor and the subscript *ind* implies that the background rate is estimated independently of the data seen. If the background rate is estimated *from data*, then the situation is more complicated. In this case, the background rate is usually estimated from sideband, i.e., a region which is located near the signal region and which contains no signal events:

$$b = \zeta_{sb} n_{sb} \ . \tag{23}$$

Here n_{sb} is the number of events observed in the sideband, and ζ_{sb} is the sideband-to-signal scale factor. Under the assumption of flat background, the scale factor is given by:

$$\zeta_{sb} = \frac{A_{sig}}{A_{sb}} \ , \tag{24}$$

where A_{sig} is the area of the signal region, and A_{sb} is the area of the sideband.

The definition (23) of the background rate is, in my opinion, incorrect. It works only if the sideband is much larger than the signal region, and thus the estimate of the expected background rate b is accurate enough. However, in a situation when the areas of signal and sideband regions are comparable, Eqn. (23), combined with Eqn. (21), overestimates the significance of a signal for large numbers of observed events $n > b$, increasing the probability of a "discovery". This is caused by the fact that the formulas (21) and (23) answer the wrong question. The correct question is: what is the probability of observing n or a larger number of events in the signal region *if the true signal rate is zero*? To answer this question, one has to assume that *all* events, that were observed in the experiment, came from background, and therefore one has to redefine:

$$b = \zeta N \ , \tag{25}$$

where N is the number of events observed in the entire region, which includes the signal region, sideband and, perhaps, an intermediate region between the signal and sideband regions, and ζ is the corresponding scale factor. Under

the assumption of flat background, the scale factor is given by:

$$\zeta = \frac{A_{sig}}{A} \ , \tag{26}$$

where A is the area of the entire region. In the same spirit, statistical significance, as defined via likelihood, is given by:

$$\sigma = \sqrt{-2\, ln\, L(0)/L_{max}} \ , \tag{27}$$

i.e., it gives a number of standard deviations from the observed signal rate, which maximizes $L(s)$, to the zero signal rate, under the assumption that $-2\, ln\, L(0)/L_{max}$ is distributed as χ^2_1. Formula (25) should be used only to estimate statistical significance of a signal; for upper limit calculation one is not allowed to assume that all observed events come from background and one has to apply the standard definition(23) of an expected background rate.

A numerical discrepancy between the definitions (23) and (25) can be illustrated on the following hypothetical example. Let us assume that the background spatial pdf is flat in the vicinity of the signal region, that the area of the sideband is equal to that of the signal region, and that one event is observed in the signal region and no events are observed in the sideband. Then an experimentalist, who uses the definition (23) of the statistical significance, would claim that she/he observes a very clean signal ($\sigma = +\infty$), while an experimentalist, who uses the definition (25), will estimate the statistical significance of the observed signal as $\sigma = 0.27$. The latter reflects the fact that, due to the specific choice of the sideband and signal regions, the observed statistic (one event) is insufficient for positive identification of the signal.

This example is purely hypothetical. Of course, in the situation when only one event is observed and the area of the sideband is comparable to that of the signal region, most of experimentalists would choose to quote an upper limit instead of a measurement. A more realistic situation occurs when the sideband is somewhat larger than the signal region, there are some events both in the signal region and in the sideband and the calculated statistical significance is close to three. Then one has to make a binary decision: if the statistical significance is larger than three, then a measurement is quoted, otherwise an upper limit is quoted. In this situation the choice of a specific procedure becomes fairly important, and the example above shows that the numerical discrepancy between the two approaches can be significant.

The conclusion that I would like to reach in this Section is that there are two situations that should be treated differently. The first situation occurs when the background rate is estimated independently of the data seen. In this

case one has to adopt the definition (22) of the background rate and therefore define statistical significance as:

$$\frac{1}{\sqrt{2\pi}} \int_{\sigma}^{\infty} e^{-x^2/2} dx = \sum_{k=n}^{\infty} e^{-\zeta_{ind} n_{ind}} \frac{(\zeta_{ind} n_{ind})^k}{k!} \ . \tag{28}$$

The second situation takes place when the background rate is estimated from the same data sample which is used to draw inference about the unknown signal rate. As shown above, in this case one has to adopt the definition (25). Statistical significance is then defined as:

$$\frac{1}{\sqrt{2\pi}} \int_{\sigma}^{\infty} e^{-x^2/2} dx = \sum_{k=n}^{\infty} e^{-\zeta N} \frac{(\zeta N)^k}{k!} \ . \tag{29}$$

These two formulas will be used to estimate upper limits, based on the statistical significance of a signal.

3.2 Upper Limit Calculation

As shown in Section 3.1, the statistical significance σ of a signal is defined either by Eqn. (28) or by Eqn. (29). I will take Eqn. (29) as an example and proceed to derive a *cumulative density function* (cdf) $P(\sigma \leq \sigma')$. If Eqn. (28) is chosen, the derivation goes through similar steps; thus, only the final result will be quoted.

When the background and signal rates are estimated from the same data sample, one can rewrite Eqn. (29) as:

$$\frac{1}{\sqrt{2\pi}} \int_{\sigma}^{\infty} e^{-x^2/2} dx = \sum_{k=n}^{\infty} e^{-\zeta(n+n_{out})} \frac{[\zeta(n+n_{out})]^k}{k!} \ , \tag{30}$$

where n is the number of events observed in the signal region, n_{out} is the number of events observed in the outer region, which covers the sideband and, perhaps, an intermediate region between the signal region and the sideband, and ζ is the corresponding scale factor. Under the assumption of flat background, the scale factor is given by:

$$\zeta = \frac{A_{sig}}{A_{sig} + A_{out}} \ , \tag{31}$$

where A_{sig} and A_{out} are the areas of the signal and outer regions respectively.

The random variables n and n_{out} are drawn from independent Poisson distributions:

$$n \sim \text{Poisson}(s+b); \quad n_{out} \sim \text{Poisson}(\lambda_{out}); \tag{32}$$

where λ_{out} is the mean of the Poisson distribution that controls the number of events in the outer region A_{out}. The best unbiased estimator of λ_{out} is the number of events actually observed outside the signal region:

$$\hat{\lambda}_{out} = n_{out,obs} \ , \tag{33}$$

and thus the equality $\lambda_{out} = n_{out,obs}$ is implied in the further discussion. In Eqn. (32) the expected number of background events b in the signal region is estimated in the traditional way: $b = \zeta_{sb} n_{sb,obs}$, and its value is not used to determine the statistical significance of the signal.

The cdf $P(\sigma \leq \sigma')$ cannot be expressed in a convenient analytical form. However, the problem can be simplified by introducing a new variable:

$$p = \sum_{k=0}^{n-1} e^{-\zeta(n+n_{out})} \frac{[\zeta(n+n_{out})]^k}{k!} \ . \tag{34}$$

Hence,

$$\frac{1}{\sqrt{2\pi}} \int_{\sigma}^{\infty} e^{-x^2/2} dx = 1 - p \ ; \quad 0 \leq p \leq 1 \ .$$

The variable p is a monotone function of σ. Therefore, its cdf can be obtained by a one-to-one transformation:

$$P(p(\sigma) \leq p(\sigma')) = P(\sigma \leq \sigma') \ , \tag{35}$$

and one can use the variable p instead of σ to set an upper limit. The cdf of p is given by:

$$1 - P(p \leq p'|s) = \sum_{n=1}^{\infty} \sum_{n_{out}=0}^{n'_{out}-1} e^{-(s+b)} \frac{(s+b)^n}{n!} e^{-\lambda_{out}} \frac{\lambda_{out}^{n_{out}}}{n_{out}!} \ , \tag{36}$$

where $n'_{out} = n'_{out}(n)$ is the smallest non-negative integer which satisfies the inequality:

$$\sum_{k=0}^{n-1} e^{-\zeta(n+n'_{out})} \frac{[\zeta(n+n'_{out})]^k}{k!} \leq p' \tag{37}$$

for the given value of n.

The infinite sum over n in Eqn. (36) converges quickly and can be easily calculated numerically. In a simulation, described in Section 4, the summation over n from 1 to 1000 was enough to achieve an accuracy of 10^{-6} or better for $(1 - P(p \leq p'|s))$.

Now one can choose the approach one would like to use. All the techniques described in Sections 2.1 and 2.2 are applicable, but now instead of a conditional pdf $f(n|s)$ one would have to use the conditional cdf $P(p \leq p'|s)$. The easiest thing to try is to employ the classical approach and define the confidence interval as $(0, s_0)$. In this case, through the reasoning described in Section 2.2, one arrives at the formula similar to Eqn. (12):

$$1 - \alpha = 1 - P(p \leq p_{obs}|s_0) \ , \tag{38}$$

where

$$p_{obs} = \sum_{k=0}^{n_{obs}-1} e^{-\zeta(n_{obs}+n_{out,obs})} \frac{\left[\zeta(n_{obs} + n_{out,obs})\right]^k}{k!} \tag{39}$$

is the observed value of the variable p.

When no events are observed in the signal region ($n_{obs} = 0$), one obtains $p_{obs} = 0$, therefore $n'_{out} = +\infty$ for any $n \geq 1$, and Eqn. (38) is reduced to:

$$1 - \alpha = 1 - e^{-(s_0+b)} \ ,$$

which coincides with the classical expression (12). This reflects the fact that in this case background fluctuations are irrelevant as the statistical significance σ is always equal to $-\infty$, no matter how many background events we observe in the sideband.

When no events are observed outside of the signal region ($\lambda_{out} = n_{out,obs} = 0$ and therefore $b = 0$), one obtains $n'_{out} \geq 1$ for any $n \geq 1$ and $p_{obs} > 0$, therefore the sum over n_{out} in Eqn. (36) is always equal to 1, and the formula (38) is reduced to:

$$1 - \alpha = 1 - e^{-s_0} \ ,$$

which again is identical to the classical expression (12). This reflects the fact that no background fluctuations are expected in this case, and hence, the statistical significance σ is a function of the signal rate s only.

If the background rate is estimated independently of the data seen, then one should start from the definition (28) of statistical significance and repeat the

same logical steps to arrive at an equation similar to Eqn. (36). The random variables n and n_{ind} are drawn from independent Poisson distributions:

$$n \sim \text{Poisson}(s+b); \quad n_{ind} \sim \text{Poisson}(\lambda_{ind}); \tag{40}$$

where λ_{ind} is the mean of the Poisson distribution that controls the number of events in an independent sample which is used to estimate the background rate. The best unbiased estimator of λ_{ind} is the number of events actually observed in the independent data sample:

$$\hat{\lambda}_{ind} = n_{ind,obs} \ , \tag{41}$$

and thus the equality $\lambda_{ind} = n_{ind,obs}$ is implied. Now the background rate used in the definition of statistical significance and the background rate used to calculate an upper limit are estimated similarly. Without losing generality, one can stipulate that:

$$\lambda_{ind} = b/\zeta_{ind} \ . \tag{42}$$

The cdf of p is now given by:

$$1 - P(p \leq p'|s) = \sum_{n=1}^{\infty} \sum_{n_{ind}=0}^{n'_{ind}-1} e^{-(s+b)} \frac{(s+b)^n}{n!} e^{-b/\zeta_{ind}} \frac{(b/\zeta_{ind})^{n_{ind}}}{n_{ind}!} \ , \tag{43}$$

where $n'_{ind} = n'_{ind}(n)$ is the smallest non-negative integer which satisfies the inequality

$$\sum_{k=0}^{n-1} e^{-\zeta_{ind} n'_{ind}} \frac{(\zeta_{ind} n'_{ind})^k}{k!} \leq p' \tag{44}$$

for the given value of n.

4 Comparison of Various Approaches Using a Toy Model

The performance of all approaches discussed in this paper is compared using the following toy model. The total observation region is defined as an interval $(-10, 10)$. The signal spatial pdf is taken to be a Gaussian with zero mean and unit variance, and the background spatial pdf is taken to be flat. Under these conditions, the signal region is defined as $(-2.5, 2.5)$, and the sidebands are defined as $(-10, -5)$ and $(5, 10)$.

The value of the expected background rate is taken to be $b = 0, 1, ..., 5$ consecutively, which corresponds to the number of events observed in the sidebands $n_{sb} = 0, 2, ..., 10$. The number of events in the intermediate region $(-5, -2.5)$ and $(2.5, 5)$ is set equal to the expected number b of background events in the signal region, since the area of the signal region is equal to that of the intermediate region. Positions of events in the sideband and intermediate regions are generated under the assumption of a uniform spatial pdf B_i. For every value of the expected background rate b, upper limits obtained with various approaches are estimated for the number n of events in the signal region varying in integer steps from 0 to 6. The results are plotted in Fig. 1. Inside the signal region all events are positioned precisely at zero. It is assumed that the background rate in this experiment is estimated from sidebands; thus, to implement the classical method of Section 3.2, formulas (38) and (39) are used. To estimate a 90% CL upper limit by the Monte Carlo technique of Section 2.3, I find the value s_0 of an upper limit by using the method of binary division. For every assumed value of s_0, I generate a Monte Carlo sample consisting of 50,000 experiments and use the formula (20) to estimate the corresponding confidence level. This procedure is repeated until the value of s_0 that gives a 90% confidence level is obtained. The required accuracy for the computation of s_0 is taken 10^{-2}. For comparison, lower and upper bounds of 90% confidence intervals obtained by the procedure [17] are shown with crosses. One cross for a specific value of n corresponds to the situation when the lower bound is strictly zero.

As shown in Fig. 1, at $n = 0$ the Bayesian method (9) with a flat prior pdf and the integration of the likelihood, described in Section 2.3, always give 2.30, while all the other approaches give somewhat smaller values. For this particular model, the likelihood integration technique of Section 2.3 turns out to be the most conservative approach, except for $b = 0$. The Bayesian method with a flat prior pdf always gives larger upper limit values than the Bayesian methods with the $1/\sqrt{s+b}$ and $1/(s+b)$ priors and both classical approaches discussed in Sections 2.2 and 3.2. The results produced by the Bayesian method with the $1/(s+b)$ prior can be obtained by shifting the corresponding results obtained with a flat prior one step to the right, e.g., $\{n = 0, m = 0\}$ and $\{n = 1, m = 1\}$ obviously produce the same result when substituted into the gamma-function $\Gamma(n - m + 1, s_0 + b)$. The classical approaches of Sections 2.2 and 3.2 fail to set an upper limit at $n \ll b$, i.e., when background dominates over signal. The classical procedure, based on the statistical significance of a signal, always gives a smaller upper limit value, compared to the standard classical approach of Section 2.2, except at $n = 0$. In the specific situation, when the expected background is zero, the Bayesian method with a flat prior pdf, both classical approaches and both likelihood techniques of Section 2.3 give identical results.

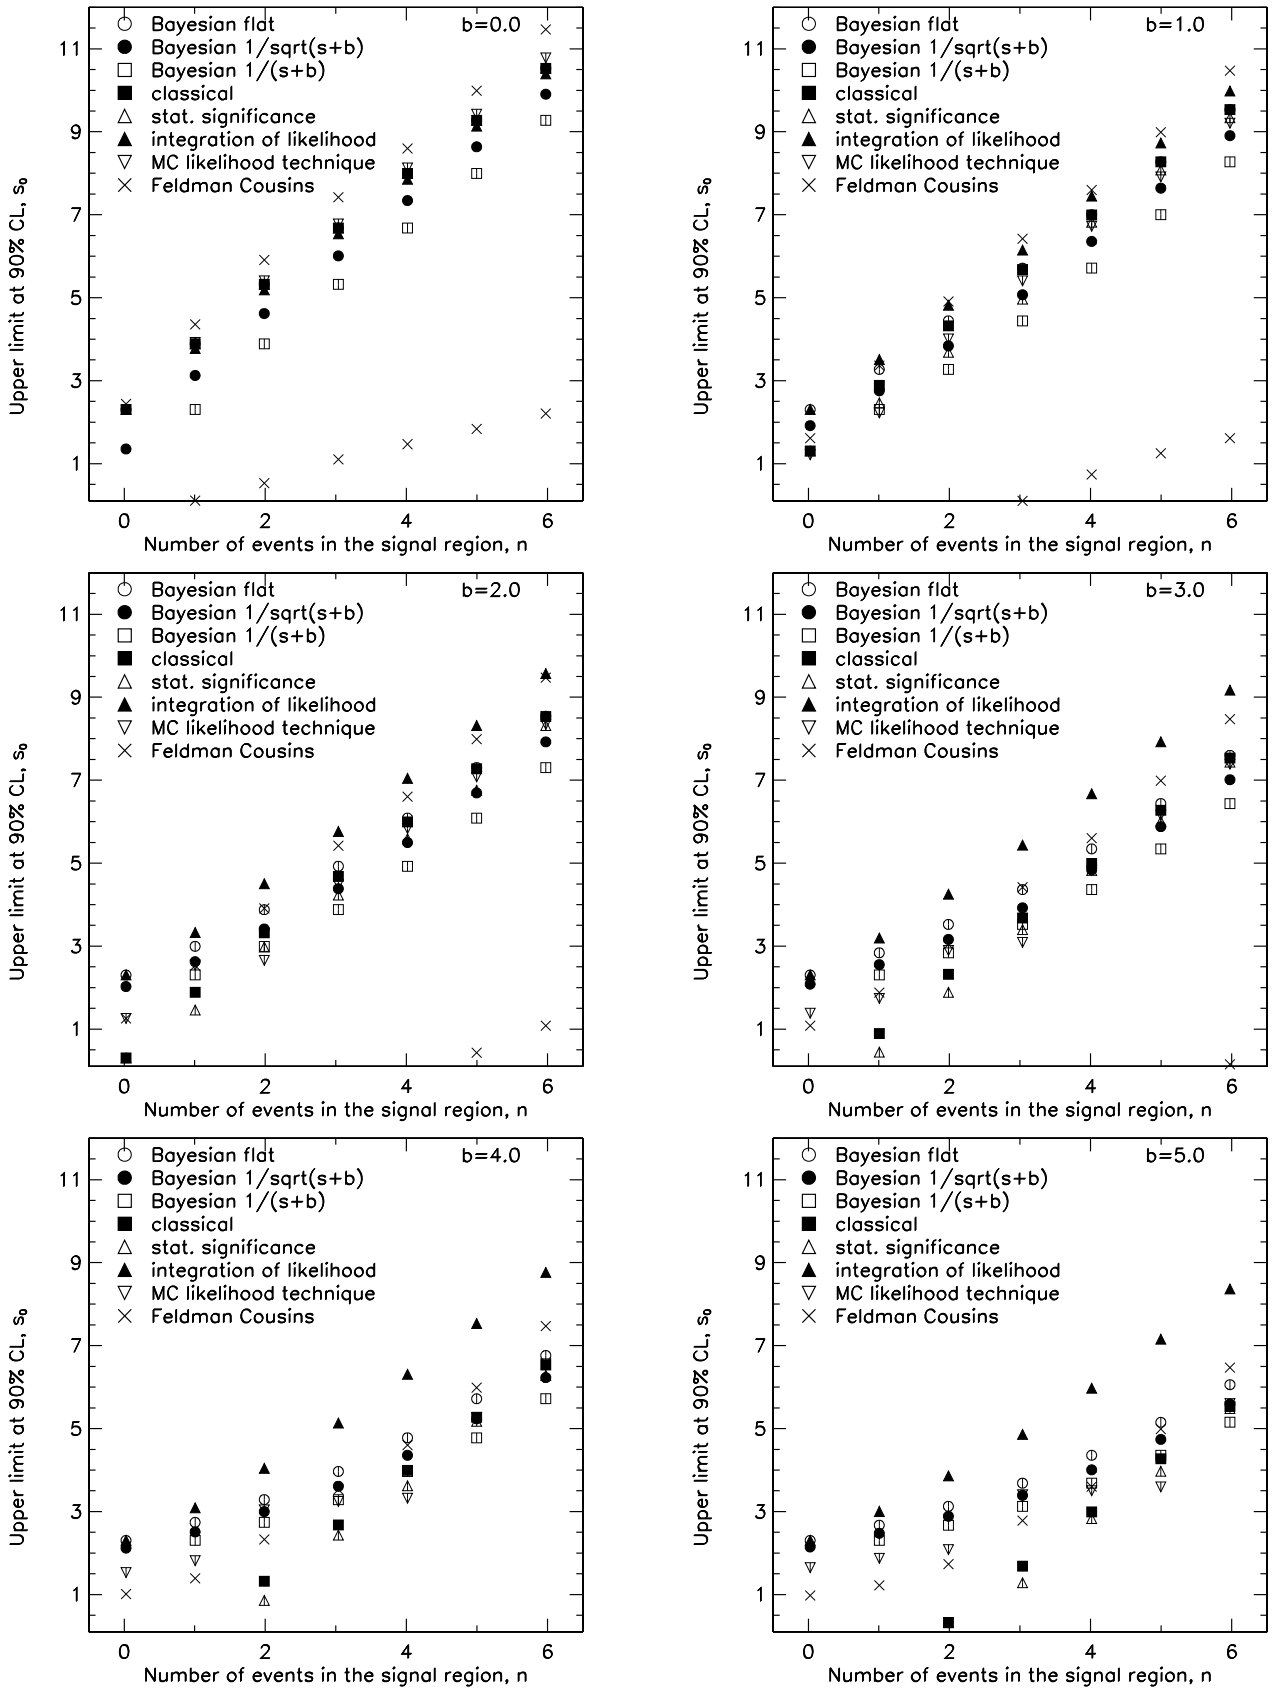

Fig. 1. Upper limits as functions of the number n of observed events in the signal region under the assumptions described in the text. The expected number of background events in the signal region takes values of 1) $b = 0$; 2) $b = 1$; 3) $b = 2$; 4) $b = 3$; 5) $b = 4$; 6) $b = 5$.

5 An Example of a CLEO Analysis: $\tau \to \mu\gamma$

In 1996 the CLEO Collaboration searched [25] for the neutrinoless decay $\tau \to \mu\gamma$ and set an upper limit, which is, so far, the most stringent limit on the $\tau \to \mu\gamma$ branching fraction. In this analysis 3 events were observed in the signal region and the expected number of background events was estimated to be 5.5. The signal Monte Carlo and data energy-vs-mass distributions are shown in Fig. 2. The Bayesian approach with a flat prior pdf was used and the upper limit value was estimated as 3.6 at 90% confidence level. This value was divided by the integrated luminosity and efficiency factor and thus an upper limit of 3.0×10^{-6} at 90% confidence level was obtained for the $\tau \to \mu\gamma$ branching fraction. In Table 1 are shown the values of upper limits for this analysis calculated with the alternative techniques.

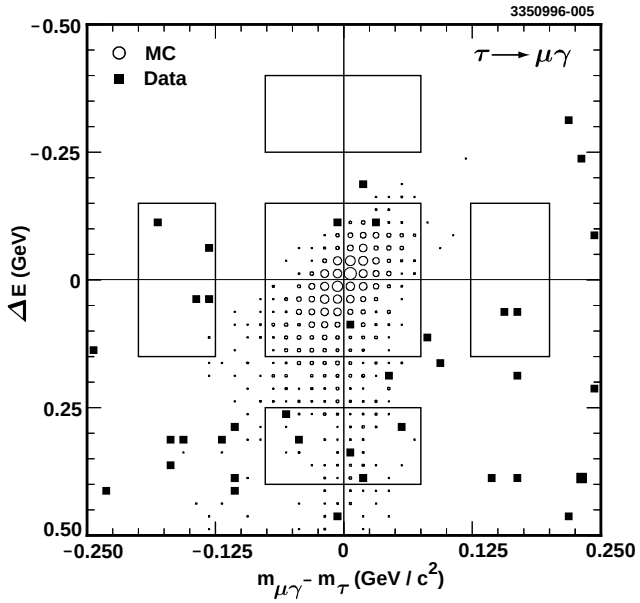

Fig. 2. $(E_{\mu\gamma} - E_{beam})$ vs $(m_{\mu\gamma} - m_\tau)$ distribution. Solid squares represent the data, open circles represent the signal Monte Carlo distribution.

The value given by the procedure [17] of Feldman and Cousins is a rough estimate only, since the combination of input parameters $\{n = 3, b = 5.5\}$ is not shown in their tables.

To implement the maximum likelihood approach of Section 2.3, the signal Monte Carlo two-dimensional distribution on the energy-vs-mass plane was fitted to a bivariate Gaussian plus a non-Gaussian tail produced by initial and final state radiation and other effects:

Table 1

Upper limits at 90% CL for the $\tau \to \mu\gamma$ analysis of Ref. [25].

Method	**Upper limit at 90% CL**
Bayesian with flat prior	**3.57**
Bayesian $1/\sqrt{s+b}$	3.30
Bayesian $1/(s+b)$	3.06
classical	1.18
classical, based on statistical significance	1.03
integration of likelihood	2.30
Monte Carlo likelihood technique	1.37
Feldman & Cousins [17]	~ 2.5

$$S_i(m,E) = \frac{1}{A+B}\left\{\frac{A}{2\pi\sigma_m\sigma_E\sqrt{1-\rho^2}} \times \right.$$
$$\times exp\left\{-\frac{1}{2(1-\rho^2)}\left[\left(\frac{m-m_0}{\sigma_m}\right)^2 - 2\rho\left(\frac{m-m_0}{\sigma_m}\right)\left(\frac{E-E_0}{\sigma_E}\right) + \left(\frac{E-E_0}{\sigma_E}\right)^2\right]\right\} +$$
$$\left. +B\epsilon(m,E)\right\} \ ; \qquad (45)$$

$$\epsilon(m,E) = \begin{cases} \frac{1}{\sqrt{2\pi}\sigma_m} exp\left[-\frac{1}{2}\left(\frac{m-m_0}{\sigma_m}\right)^2\right] \frac{1}{\sigma_E\Gamma(\alpha)\beta^\alpha}\left(\frac{E_0-E}{\sigma_E}\right)^{\alpha-1} exp\left[-\frac{E_0-E}{\beta\sigma_E}\right] & \text{if } E < E_0 \\ 0 & \text{otherwise} \end{cases}$$

where A, B, σ_m, σ_E, ρ, $m_0(\approx m_\tau)$, $E_0(\approx E_{beam})$, α and β are the fit parameters. The background spatial pdf B_i was obtained by fitting data events observed in the vicinity of the signal region to a linear function:

$$B_i(m,E) = \frac{1}{m_2-m_1}\frac{1}{(a_0-a_1E')(E_2-E_1)+0.5a_1(E_2^2-E_1^2)}\left[a_0+a_1(E-E')\right] \ , \qquad (46)$$

where a_0, a_1 and E' are the fit parameters, and (m_1, m_2) and (E_1, E_2) are the limits defining the fit region. The three events observed in the signal region are located far from the peak of the signal Monte Carlo distribution. Thus, the maximum likelihood fit treats these events as background, and the extracted signal rate is consistent with zero. This explains why the likelihood integration technique gives the value of 2.30.

6 Conclusion

There is no such thing as the "best" procedure for upper limit estimation. An experimentalist is free to choose any procedure she/he likes, based on her/his belief and experience. The only requirement is that the chosen procedure must have a strict mathematical foundation. The Bayesian method with a flat prior and the likelihood integration technique of Section 2.3 seem to have been the two most popular choices in the past few years. Typically, these two approaches give the most conservative values of upper limits. The purpose of this note is to show that there are other approaches, equally justified by mathematical formalism, which produce less conservative estimates.

7 Acknowledgements

The author thanks Prof. W.T. Ford, Prof. R. Stroynowski, Prof. H.L. Gray and Dr. I. Volobouev for useful comments and suggestions.

References

[1] T. Bayes, *Facsimiles of Two Papers by Bayes* (Hafner Publ. Co., 1963).

[2] P.S. Laplace, *Philosophical Essay on Probabilities* (Dover Publications, 1951).

[3] H. Jeffreys, *Theory of Probability* (3rd ed., Oxford at Clarendon, 1961).

[4] H. Jeffreys, *Proc. of the Royal Soc. of London* **A 196** (1946) 453.

[5] E.T. Jaynes, *IEEE Trans. on Systems Sci. and Cyber.* **SSC-4, 3** (1968) 227.

[6] H.B. Prosper, *Phys. Rev.* **D 37, N 5** (1988) 1153.

[7] G.E.P. Box and G.C. Tiao, *Bayesian Inference in Statistical Analysis* (Wiley Classics, 1992).

[8] O. Helene, *Nucl. Instr. and Meth.* **212** (1983) 319.

[9] O. Helene, *Nucl. Instr. and Meth. in Phys. Res.* **228** (1984) 120.

[10] B. Escoubes, S. De Unamuno and O. Helene, *Nucl. Instr. and Meth. in Phys. Res.* **A 257** (1987) 346.

[11] O. Helene, *Nucl. Instr. and Meth. in Phys. Res.* **A 300** (1991) 132.

[12] G. Zech, *Nucl. Instr. and Meth. in Phys. Res.* **A 277** (1989) 608.

[13] R.M. Barnett et al. (Particle Data Group), *Phys. Rev.* **D 54** (1996) 166.

[14] C. Caso et al. (Particle Data Group), *The European Physical Journal* **C 3** (1998) 176.

[15] J. Neyman, *Philos. Trans. of the Royal Soc. of London* **A 236** (1937) 333.

[16] F. Garwood, *Biometrika* **28** (1936) 437.

[17] G.J. Feldman and R.D. Cousins, *Phys. Rev.* **D 57** (1998) 3873.

[18] E.L. Crow and R.S. Gardner, *Biometrika* **46** (1959) 441.

[19] R.D. Cousins, *Am. J. of Phys.* **63 (5)** (1995) 398.

[20] W.T. Eadie et al., *Statistical Methods in Experimental Physics* (North-Holland, Amsterdam, 1971).

[21] L. Lyons, W. Allison and J. Panella Comellas, *Nucl. Instr. and Meth.* **A 245** (1986) 530.

[22] R. Barlow, *Nucl. Instr. and Meth.* **A 297** (1990) 496.

[23] B. Behrens et al. (CLEO), *Phys. Rev. Lett.* **80** (1998) 3710.

[24] T. Bergfeld et al. (CLEO), *Phys. Rev. Lett.* **81** (1998) 272.

[25] K.W. Edwards et al. (CLEO), *Phys. Rev.* **D 55**, (1997) 3919.